

Lattice template placement for coherent all-sky searches for gravitational-wave pulsars

Karl Wette*

Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut), Callinstr. 38, 30167 Hannover, Germany

(Dated: 2014-12-22 09:50:44 +0100; commitID: cb5ec5f-CLEAN)

All-sky, broadband, coherent searches for gravitational-wave pulsars are restricted by limited computational resources. Minimizing the number of templates required to cover the search parameter space, of sky position and frequency evolution, is one important way to reduce the computational cost of a search. We demonstrate a practical algorithm which, for the first time, achieves template placement with a minimal number of templates for an all-sky search, using the reduced supersky parameter-space metric of Wette and Prix [Phys. Rev. D **88**, 123005 (2013)]. The metric prescribes a constant template density in the signal parameters, which permits that templates be placed at the vertices of a lattice. We demonstrate how to ensure complete coverage of the parameter space, including in particular at its boundaries. The number of templates generated by the algorithm is compared to theoretical estimates, and to previous predictions by Brady *et al.* [Phys. Rev. D **57**, 2101 (1998)]. The algorithm may be applied to any search parameter space with a constant template density, which includes semicoherent searches and searches targeting known low-mass X-ray binaries.

PACS numbers: 04.80.Nn, 95.55.Ym, 95.75.Pq, 97.60.Jd

I. INTRODUCTION

Gravitational-wave pulsars are rapidly-spinning neutron stars which could be emitting gravitational radiation if nonaxisymmetrically deformed, due to various mechanisms which may support such a deformation (see [1, 2] for reviews of emission mechanisms, and [3, 4] for predictions of the maximum supportable nonaxisymmetry). Their detection by ground-based interferometric detectors with kilometer-long arms, such as LIGO [5] and Virgo [6], is one of the great challenges of gravitational-wave physics. Searches for gravitational-wave pulsars in data from the first generation of interferometric detectors [e.g. 7–9] have, to date, not yielded a detection. Indeed, gravitational-wave pulsars are expected to be difficult to detect even by the next generation of detectors [10–12], which are currently under construction. Maximizing the chance of a detection requires investment in highly-optimized data analysis techniques, and large-scale computing resources such as Einstein@Home [8].

The signals emitted by gravitational-wave pulsars are characteristically continuous, narrow-band, and quasi-sinusoidal, and are believed to be well-modeled by a parameterized template waveform family. The most sensitive search method is thus coherent matched filtering against a bank of templates, whose parameters are chosen from a space of interest; this yields the detection statistic commonly known in the field as the \mathcal{F} -statistic [13, 14]. This method is, however, too computationally intensive to search year-long data sets and wide parameter spaces, such as all-sky broadband-frequency searches for undiscovered gravitational-wave pulsars. For such searches, a semicoherent search is employed: the data are divided into shorter segments, each of which is coherently

matched filtered, and the results from each segment combined using a computationally cheaper, less sensitive incoherent method [e.g. 15–18]. The most sensitive semicoherent search setup, as a function of e.g. the number and length of segments, has been studied in [19, 20].

An important component of a wide-parameter-space \mathcal{F} -statistic search for gravitational-wave pulsars is the *metric*, or distance function, associated with the parameter space [15, 21]. It quantifies how far apart templates may be separated such that any signal in the search parameter space will be recovered with a prescribed maximum *mismatch*, or fractional loss in squared signal-to-noise ratio. If the metric is not itself a function of the template parameters, the number of templates required to cover the parameter space, and hence the computational cost of the search, can be straightforwardly minimized using the theory of lattices and sphere coverings [22, 23]. This is the case for the parameters which describe the gravitational-wave frequency evolution of a pulsar: its frequency at a given reference time, and its frequency time derivatives, or *spindowns*. It is more difficult, however, to parameterize the pulsar’s sky position such that this is true.

Recently, in Wette and Prix [24, hereafter Paper I] we proposed a choice of sky and frequency parameters, with respect to which a close approximation to the metric, the *reduced supersky metric*, is constant. With respect to previously proposed approximations [18, 21, 25–27], the new metric has no restrictions on the time span of data which can be coherently analyzed, and is numerically well-conditioned which eases its practical use.

Building on the work in Paper I, this paper presents further investigations of the reduced supersky metric, with a primary focus on its practical use for template placement. Section II reviews relevant background information. In Section III, we refine numerical simulations, used in Paper I to test the reduced supersky metric, to more accurately reflect its use in a real search. In Sec-

* karl.wette@aei.mpg.de

tion IV, we demonstrate that the simulated mismatch predictions of the reduced supersky metric can be realized in a practical search using lattice template placement; this is the first time this has been demonstrated for an all-sky search. In Section V, we investigate an important property of the reduced supersky metric: the number of templates it predicts are needed to cover the parameter-space of a coherent all-sky search. Directions for future work are discussed in Section VI.

II. BACKGROUND

This section briefly reviews background information relevant to this paper. We review the gravitational-wave pulsar signal model and parameter-space metric in Section II A, and the supersky and reduced supersky metrics in Section II B; see Paper I for further details. We briefly introduce lattices in Section II C; see e.g. [22] for a comprehensive treatment.

A. The signal model and parameter-space metric

The signal model of a gravitational-wave pulsar [13], $h(t, \mathcal{A}^i, \vec{\lambda})$, is a function of four amplitude parameters \mathcal{A}^i , and a number of phase evolution parameters $\vec{\lambda}$. The \mathcal{A}^i are functions of the gravitational-wave strain amplitude h_0 and initial phase ϕ_0 , and the pulsar's angles of inclination ι and polarization ψ . The $\vec{\lambda}$ are the pulsar's sky position, represented by a unit vector \vec{n} pointing from the Solar System barycenter (SSB) to the pulsar, frequency $f \equiv f^{(0)}$ at SSB reference time t_0 , and s_{\max} spindowns $f^{(s)} \equiv d^s f / dt^s|_{t=t_0}$. The signal model can be written as $h(t, \mathcal{A}, \vec{\lambda}) = \sum_{i=1}^4 \mathcal{A}^i h_i(t, \vec{\lambda})$, where the $h_i(t, \vec{\lambda})$ are four time- and phase-parameter-dependent functions. We restrict our attention to *isolated* gravitational-wave pulsars, i.e. those without a binary companion (but see the discussion in Section VI).

The \mathcal{F} -statistic [13, 14] matched-filters data from a gravitational-wave detector against the signal model, and further maximizes over the unknown amplitudes \mathcal{A}^i ; it is therefore a function of only the phase parameters $\vec{\lambda}$. An \mathcal{F} -statistic search computes $\mathcal{F}(\vec{\lambda})$ for a discrete set of values $\{\vec{\lambda}^t\}$, which comprise the template bank. It is unlikely, however, that the parameters $\vec{\lambda}^s$ of any signal in the data will precisely match one of the $\vec{\lambda}^t$ in the template bank. Therefore, the signal will be recovered with a signal-to-noise ratio $\rho(\mathcal{A}, \vec{\lambda}^s; \vec{\lambda}^t)$ lower than for a perfect match $\rho(\mathcal{A}, \vec{\lambda}^s; \vec{\lambda}^s)$. The mismatch μ_0 is defined to be [21, 24]

$$\mu_0 = \frac{\rho^2(\mathcal{A}, \vec{\lambda}^s; \vec{\lambda}^s) - \rho^2(\mathcal{A}, \vec{\lambda}^s; \vec{\lambda}^t)}{\rho^2(\mathcal{A}, \vec{\lambda}^s; \vec{\lambda}^s)}. \quad (1)$$

For small differences $\Delta \vec{\lambda} = \vec{\lambda}^s - \vec{\lambda}^t$, a second-order Taylor

expansion of Eq. (1) yields the metric \mathbf{g} :

$$\mu_0 \approx \Delta \vec{\lambda}^T \frac{-1}{2\rho^2(\mathcal{A}, \vec{\lambda}^s; \vec{\lambda}^s)} \frac{\partial \rho^2(\mathcal{A}, \vec{\lambda}^s; \vec{\lambda})}{\partial \vec{\lambda}} \bigg|_{\vec{\lambda}=\vec{\lambda}^s} \Delta \vec{\lambda} \quad (2)$$

$$= \Delta \vec{\lambda}^T \mathbf{g} \Delta \vec{\lambda}, \quad (3)$$

where \cdot^T denotes matrix transposition and transformation between row and column vectors.

The elements of \mathbf{g} are complicated functions of both amplitude and phase parameters [21]. A useful approximation which depends only on the phase parameters is the *phase metric* \mathbf{g}_ϕ [15, 21], with elements

$$[\mathbf{g}_\phi]_{ij} = \left\langle \frac{\partial \phi(t, \vec{\lambda})}{\partial \lambda_i} \frac{\partial \phi(t, \vec{\lambda})}{\partial \lambda_j} \right\rangle - \left\langle \frac{\partial \phi(t, \vec{\lambda})}{\partial \lambda_i} \right\rangle \left\langle \frac{\partial \phi(t, \vec{\lambda})}{\partial \lambda_j} \right\rangle, \quad (4)$$

where $\langle x(t) \rangle = \int_{t_{\text{start}}}^{t_{\text{start}}+T} dt x(t)/T$, t_{start} is the start time and T is the time span of the data segment being searched. The function $\phi(t, \vec{\lambda})$ denotes the phase of the gravitational-wave pulsar signal in a given detector at time t , and is approximately

$$\frac{\phi(t, \vec{\lambda})}{2\pi} \approx \sum_{s=0}^{s_{\max}} f^{(s)} \frac{(t - t_0)^{s+1}}{(s+1)!} + \frac{\vec{r}(t) \cdot \vec{n}}{c} f_{\max}, \quad (5)$$

where $\vec{r}(t)$ is the detector position relative to the SSB, and f_{\max} is a constant usually chosen conservatively to be the maximum of the instantaneous frequency $f(t)$ over T ; see Paper I. If, as is the case for $f^{(s)}$, the phase $\phi(t, \vec{\lambda})$ is linear in a parameter λ_i , then by Eq. (4) \mathbf{g}_ϕ is independent of λ_i .

B. The supersky and reduced supersky metrics

Paper I proposed adopting the three components of \vec{n} as sky position parameters; it follows that the phase metric in the coordinates $(\vec{n}, f^{(s)})$, the *supersky* metric \mathbf{g}_{ss} , is constant. A caveat is that this choice of parameters embeds the two-dimensional space of possible sky positions, represented by the two-sphere $|\vec{n}| = 1$, in the three-dimensional space of possible vectors $\vec{n} \in \mathbb{R}^3$. Since the physically interesting parameter space is now only a subspace of the full parameter space, template placement is no longer straightforward.

To reduce the dimensionality of \mathbf{g}_{ss} , Paper I outlined a procedure which selects a two-dimensional subspace in \mathbb{R}^3 such that the metric in this subspace, the *reduced supersky* metric \mathbf{g}_{rss} , is a close approximation to \mathbf{g}_{ss} . The selected subspace is the plane perpendicular to the eigenvector corresponding to the smallest eigenvalue of $\mathbf{g}_{\text{ss}, \vec{n}\vec{n}}$, the 3×3 block of \mathbf{g}_{ss} pertaining only to the sky parameters. Along this axis, the supersky mismatch μ_{ss} changes slowest as a function of differences in sky position; dropping this dimension therefore introduces the smallest possible error in an approximation to μ_{ss} .

In order to reliably perform the above procedure, however, the numerical ill-conditionedness of \mathbf{g}_{ss} must be addressed. It was found that the ill-conditionedness arises from the near-linear relation, for $T \ll 1$ year, between $\phi(t, \vec{\lambda})$ as a function of the orbital motion of the Earth, and as a function of the frequency evolution of the pulsar. This effect has also been observed and exploited in previous work [26, 27]. In Paper I it is used to devise a linear transformation of the coordinates $(\vec{n}, f^{(s)})$ which removes a linear fit to the component of \mathbf{g}_{ss} due to orbital motion by the component due to frequency evolution.

A further linear transformation of the supersky metric removes the correlations between \vec{n} and $f^{(s)}$, i.e. such that the sky-frequency blocks $\mathbf{g}_{\text{ss}, \vec{n} f^{(s)}}$ and $\mathbf{g}_{\text{ss}, f^{(s)} \vec{n}}$ of \mathbf{g}_{ss} are zero. The sky position is then expressed in the eigenbasis of the sky-sky block $\mathbf{g}_{\text{ss}, \vec{n} \vec{n}}$, and the dimension corresponding to the smallest eigenvalue is dropped, yielding the reduced supersky metric \mathbf{g}_{rss} . The metric is constant, and its *condition number*, the ratio of its largest to smallest eigenvalues, is of order unity. The associated coordinates are $(n_a, n_b, \nu^{(s)})$, where n_a and n_b are the sky coordinates corresponding to the two largest eigenvalues of $\mathbf{g}_{\text{ss}, \vec{n} \vec{n}}$, and the $\nu^{(s)}$ are frequency and spindown coordinates linear in $f^{(s)}$ and $\vec{n} = (n_a, n_b, \pm \sqrt{1 - n_a^2 - n_b^2})$.

C. Lattices

A *lattice* is a set of n -dimensional points $\{\vec{p}^\ell\} \subset \mathbb{R}^n$ which is closed under vector addition and subtraction, i.e. if \vec{p}_1^ℓ and \vec{p}_2^ℓ are lattice points, then so are $\vec{p}_1^\ell \pm \vec{p}_2^\ell$. It follows that a lattice can be generated by a linear transformation from integer vectors $\vec{k} \in \mathbb{Z}^n$ to lattice points $\vec{p}^\ell \in \{\vec{p}^\ell\}$, represented by an $n \times n$ *generator matrix* \mathbf{G} .

The *covering radius* R of a lattice is defined such that every point $\vec{p} \in \mathbb{R}^n$ is within a Euclidean distance R of some lattice point $\vec{p}^\ell \in \{\vec{p}^\ell\}$, i.e.

$$\min_{\vec{p}^\ell \in \{\vec{p}^\ell\}} \|\vec{p} - \vec{p}^\ell\| \leq R, \quad (6)$$

and no smaller R satisfies this inequality. The region of \mathbb{R}^n defined by Eq. (6) is the *covering sphere* centered on the point \vec{p}^ℓ .

The ratio of the volume contained in a covering sphere to the volume per lattice point, the *normalized thickness* $\theta = R^n / \sqrt{\det \mathbf{G}}$, is a fundamental property of a lattice. The number of templates needed to cover a given parameter space will be minimized by the lattice with the smallest normalized thickness.

III. REFINED NUMERICAL SIMULATIONS

In Paper I the mismatch predictions of the reduced supersky were investigated using numerical simulations. These simulations generated random parameter offsets $\Delta \vec{\lambda}$, and examined the difference between the mismatch




FIG. 1. Histograms of normalized mismatch μ/μ_{max} expected from lattice placement using (a) \mathbb{Z}^n lattices and (b) A_n^* lattices, with $n = 2-5$. Each histogram was generated using 10^9 simulated points.

$\mu_{\text{rss}} = \Delta \vec{\lambda}^T \mathbf{g}_{\text{rss}} \Delta \vec{\lambda}$ predicted by the metric, and the mismatch μ_0 calculated from the \mathcal{F} -statistic using Eq. (1); see the Appendix of Paper I for details. In Sections III A, III B, and III C, we describe three refinements to the simulations performed in Paper I. The results of the new simulations are presented in Section III D.

A. Realistic mismatch distributions

In Paper I, the parameter offsets $\Delta \vec{\lambda}$ were sampled to produce a uniform distribution in μ_{rss} , up to some maximum μ_{max} . When placing templates using a lattice, however, the expected distribution of mismatches is not uniform, and depends on the geometry of the lattice. Figure 1 plots examples of the mismatch distributions [cf. 28] expected when placing templates using two lattice families: the \mathbb{Z}^n lattices, which generalize the two-dimensional square lattice to higher dimensions; and the A_n^* lattices, which generalize the two-dimensional hexagonal lattice. The A_n^* lattices have the smallest known normalized thicknesses in low dimensions [22].

The simulations presented in this paper sample $\Delta \vec{\lambda}$ to produce the distribution in μ_{rss} expected when using an A_n^* lattice to place templates in a real search. This is achieved by choosing a random point $\vec{\lambda} \in \mathcal{P}$, then finding

the lattice template point $\vec{\lambda}^t \in \{\vec{\lambda}^t\}$ that minimizes

$$\mu_{\text{rss}} = (\vec{\lambda} - \vec{\lambda}^t)^T \mathbf{g}_{\text{rss}} (\vec{\lambda} - \vec{\lambda}^t) = \Delta \vec{\lambda}^T \mathbf{g}_{\text{rss}} \Delta \vec{\lambda}. \quad (7)$$

As a consequence, $\mu_{\text{rss}}/\mu_{\text{max}}$ is distributed according to the mismatch distribution of an A_n^* lattice (Fig. 1b).

Let \mathbf{A} denote the Cholesky factorization [e.g. 29] of \mathbf{g}_{rss} , i.e. a lower triangular matrix satisfying $\mathbf{g}_{\text{rss}} = \mathbf{A}\mathbf{A}^T$. Then minimizing Eq. (7) is equivalent to minimizing $\|\vec{p} - \vec{p}^\ell\|$, where $\vec{p} = \mathbf{A}^T \vec{\lambda}$, and $\vec{p}^\ell = \mathbf{A}^T \vec{\lambda}^t$ is the point in the A_n^* lattice corresponding to the template $\vec{\lambda}^t$. For many lattices, including A_n^* , efficient algorithms exist which, given \vec{p} , find the \vec{p}^ℓ which minimizes $\|\vec{p} - \vec{p}^\ell\|$ (see Section IV C). Once \vec{p}^ℓ is found, $\Delta \vec{\lambda}$ is found via $\Delta \vec{\lambda} = [\mathbf{A}^T]^{-1}(\vec{p} - \vec{p}^\ell)$, where \cdot^{-1} denotes matrix inversion.

B. Fixed reference time

The simulations in Paper I tested the reduced supersky metric at different start times t_{start} within a 1-year period. The reference time t_0 , which enters the calculation of \mathbf{g}_{rss} via Eq. (5), was always set to $t_0 = t_{\text{start}} + 0.5T$, the mid-time of the data segment for which \mathbf{g}_{rss} was being calculated.

In this paper, we instead fix $t_0 = \text{UTC } 2007-06-30 \text{ } 00:03:06$, and perform the simulations at 25 values of $\Delta t_{\text{start}} = t_{\text{start}} - t_0$ from -180 to +180 days in steps of 15 days. This setup reflects that of a semicoherent search where, in order to combine \mathcal{F} -statistic values from several coherently-analyzed data segments together, it is convenient if the template banks of each data segment (and hence the metrics used to generate them) are defined at the same reference time t_0 . In addition, when $t_0 \neq t_{\text{start}} + 0.5T$ the frequency/spindown off-diagonal elements of the metric, $g_{\text{rss}}(\nu^{(s)}, \nu^{(s')})$ with $s \neq s'$, are non-negligibly nonzero; the effect this has on the mismatch predictions of \mathbf{g}_{rss} was not tested in Paper I.

C. Multiple detectors

Paper I tested the reduced supersky metric computed at a single detector location, that of the LIGO Hanford detector. In this paper, we average the reduced supersky metrics computed at the locations of the LIGO Hanford and Livingston detectors. This is an ad-hoc choice, as the phase metric approximation is defined for a single detector only [21]. Nevertheless, the choice proved successful, as demonstrated by the mismatch predictions of the reduced supersky metric presented in the next section.

D. Simulation results

The simulations presented here follow the procedure outlined in the Appendix of Paper I, with the refine-

ments described above. Simulations are performed at: fixed values of T , in steps of 2 days, from 1 to 31 days for first spindown and 11 to 31 days for second spindown; fixed values of Δt_{start} , as given in Section IIIB, and fixed values of f_{max} from 50 to 1000 Hz, as given in the Appendix of Paper I. Mismatches are compared using their *relative error*, defined following Paper I to be

$$\varepsilon(\mu_a, \mu_b) = \frac{\mu_a - \mu_b}{0.5(\mu_a + \mu_b)}, \quad \mu_a, \mu_b \geq 0. \quad (8)$$

Figures 2 and 3 plot the relative errors $\varepsilon(\mu_0, \mu_{\text{rss}})$ and $\varepsilon(\mu_{\text{ss}}, \mu_{\text{rss}})$ between mismatches μ_{rss} and μ_{ss} predicted by the reduced supersky and supersky metrics, \mathbf{g}_{rss} and \mathbf{g}_{ss} respectively, and the mismatch μ_0 calculated from the \mathcal{F} -statistic. The simulations presented in Fig. 2 test metrics computed at time spans T from 1 to 31 days using random offsets in the sky coordinates (n_a, n_b) , frequency ν , and first spindown $\dot{\nu}$; those presented in Fig. 3 use T from 11 to 31 days and include additional random offsets in second spindown $\ddot{\nu}$. Figures 2a and 2c are comparable to Figs. 2f and 14 of Paper I, while Figs. 3a and 3c are comparable to Figs. 17a and 17b of that paper.

The simulations presented here show comparable, even improved, mismatch predictions by the reduced supersky metric compared to Paper I, despite using a slightly larger maximum mismatch of 0.3 (cf. 0.2 in Paper I). For example, Fig. 2f of Paper I shows $\varepsilon(\mu_0, \mu_{\text{rss}}) \sim -0.5$ at $T = 1$ day, whereas Fig. 2a shows $\varepsilon(\mu_0, \mu_{\text{rss}}) \sim -0.15$ at the same T . This is expected; as seen in Fig. 1, the mismatch distribution expected from using a A_4^* lattice includes fewer large mismatches (e.g. $\mu/\mu_{\text{max}} \gtrsim 0.8$) than would a uniform distribution. We expect the metric to perform worse at larger mismatches, due to the deterioration of the metric approximation to the \mathcal{F} -statistic mismatch (see Sec. IV A and Fig. 7 of Paper I). The small relative errors seen in the bottom rows of Figs. 2 and 3 confirm the close agreement between the supersky \mathbf{g}_{ss} and reduced supersky \mathbf{g}_{rss} metrics seen in Fig. 14 of Paper I.

No deterioration in the mismatch predictions of the reduced supersky metric as a function of Δt_{start} is observed in Figs. 2b, 2d, 3b, and 3d. This indicates that the derivation of the reduced supersky metric is robust to the difference between reference time t_0 and segment start time t_{start} .

Figure 4 compares the distributions of mismatch in μ_{rss} , μ_0 , and μ_{ss} , averaged over all simulation parameters. The mismatch distributions of μ_{rss} closely resemble the desired mismatch distributions for an A_n^* lattice; see Fig. 1b with $n = 4$ (first spindown only) and $n = 5$ (both first and second spindown). Given the close agreement between the reduced supersky and supersky metrics seen in Figs. 2c, 2d, 3c, and 3d, it is expected that the mismatch distributions of μ_{ss} resemble those of μ_{rss} .

Likewise, the differences between the μ_0 and μ_{rss} mismatch distributions seen in Fig. 4 are also expected, given that the reduced supersky metric does not perfectly predict the \mathcal{F} -statistic mismatch (Figs. 2a, 2b, 3a, and 3b).




FIG. 2. Relative errors between μ_{rss} and μ_0 (top row), and between μ_{rss} and μ_{ss} (bottom row). Only first spindown is used. Plotted are the median (solid line), the 25th–75th percentile range (error bars), and the 2.5th (short-dashed line) and 97.5th (long-dashed line) percentiles of relative errors: as a function of T , averaged over Δt_{start} and f_{max} (left column); and as a function of Δt_{start} , averaged over T and f_{max} (right column). An A_n^* lattice is used to generate random parameter offsets with $\mu_{\text{rss}} \leq 0.3$.

That the μ_0 distributions peak at lower mismatches than those of μ_{rss} is consistent with the reduced supersky metric overestimating the \mathcal{F} -statistic mismatch; see the discussion of Fig. 7 in Paper I. On the other hand, the means of the distributions are very similar: for first spindown only, the mean μ_{rss} mismatch is 0.15 and the mean μ_0 mismatch is 0.14; for both first and second spindown, the means are 0.17 and 0.15 respectively.

IV. LATTICE TEMPLATE PLACEMENT

In the previous section, we confirmed that the \mathcal{F} -statistic mismatch is well-predicted by the reduced supersky metric, when the distribution of μ_{rss} resembles that expected for lattice template placement. In this section, we demonstrate that such mismatch distributions are realized by a practical implementation [30] of a lattice template bank. The implementation is based on one developed in [31], which was used in the gravitational-wave pulsar search presented in [32]; lattice template placement is also discussed in [23]. Constituent parts of the implementation are described in Sections IV A, IV B, and IV C; tests of the implementation are presented in Section IV D. An issue pertaining to the coverage of parameter-space boundaries is discussed in Section IV E.

A. Parameter-space representation

We consider an n -dimensional parameter space $\mathcal{P} \subset \mathbb{R}^n$, within which templates are points $\vec{\lambda}^t = (\lambda_0^t, \lambda_1^t, \dots, \lambda_{n-1}^t) \in \mathbb{R}^n$. We prescribe that \mathcal{P} be represented by inequalities of the form:

$$\begin{aligned} \lambda_0^{\min} &\leq \lambda_0^t \leq \lambda_0^{\max}, \\ \lambda_1^{\min}(\lambda_0^t) &\leq \lambda_1^t \leq \lambda_1^{\max}(\lambda_0^t), \\ \lambda_2^{\min}(\lambda_0^t, \lambda_1^t) &\leq \lambda_2^t \leq \lambda_2^{\max}(\lambda_0^t, \lambda_1^t), \\ &\dots, \\ \lambda_{n-1}^{\min}(\dots, \lambda_{n-2}^t) &\leq \lambda_{n-1}^t \leq \lambda_{n-1}^{\max}(\dots, \lambda_{n-2}^t). \end{aligned} \quad (9)$$

The lowest dimension λ_0^t is bounded by two constants λ_0^{\min} and λ_0^{\max} ; the dimension λ_1^t by two functions $\lambda_1^{\min}(\lambda_0^t)$ and $\lambda_1^{\max}(\lambda_0^t)$, depending only on λ_0^t ; the dimension λ_2^t by two functions $\lambda_2^{\min}(\lambda_0^t, \lambda_1^t)$ and $\lambda_2^{\max}(\lambda_0^t, \lambda_1^t)$, depending only on λ_0^t and λ_1^t ; and so on up to λ_{n-1}^t . The motivation for this representation is presented in the next section. Figure 5 shows an example of a parameter space which is described in this form.

The region of \mathcal{P} covered by a template $\vec{\lambda}^t$ is the set of points $\{\vec{\lambda}\}$ such that

$$(\vec{\lambda} - \vec{\lambda}^t)^T \mathbf{g}_{\text{rss}}(\vec{\lambda} - \vec{\lambda}^t) \leq \mu_{\text{max}}; \quad (10)$$




FIG. 3. Relative errors between μ_{rss} and μ_0 (top row), and between μ_{rss} and μ_{ss} (bottom row). Both first and second spindown are used. Plotted are the median (solid line), the 25th–75th percentile range (error bars), and the 2.5th (short-dashed line) and 97.5th (long-dashed line) percentiles of relative errors: as a function of T , averaged over Δt_{start} and f_{max} (left column); and as a function of Δt_{start} , averaged over T and f_{max} (right column). An A_n^* lattice is used to generate random parameter offsets with $\mu_{\text{rss}} \leq 0.3$.

this expression describes an n -dimensional *metric ellipse*. The *metric ellipse bounding box* of \mathbf{g}_{rss} is the smallest n -dimensional coordinate box which contains the metric ellipse; its widths β_i in each dimension $i = 0, \dots, n-1$ are the minimum required to satisfy $|\lambda_i - \lambda_i^t| \leq 0.5\beta_i$, for all points $\vec{\lambda}$ satisfying Eq. (10). The widths $\vec{\beta}$ are computed from the metric via [e.g. 31]

$$\beta_i = 2\sqrt{\mu_{\text{max}}[\mathbf{g}_{\text{rss}}^{-1}]_{ii}}, \quad (11)$$

where $[\mathbf{g}_{\text{rss}}^{-1}]_{ii}$ is the i th diagonal element of $\mathbf{g}_{\text{rss}}^{-1}$.

Care must be taken at the boundaries of \mathcal{P} to ensure that it is completely covered. Figure 6 illustrates a common situation where simply laying templates inside the given bounds of \mathcal{P} does not fully cover it. By extending the boundaries of \mathcal{P} by half of the extent of the bounding box, given by Eq. (11), complete coverage is achieved.

B. Lattice template generation

The template points $\vec{\lambda}^t$ are generated from integer vectors $\vec{k} \in \mathbb{Z}^n$ by

$$\vec{\lambda}^t = \mathbf{T}\vec{k}, \quad (12)$$

where the matrix \mathbf{T} transforms from \mathbb{Z}^n to \mathcal{P} . Each template point $\vec{\lambda}^t$ is placed at a vertex of the lattice. It

can be shown [31] that, if \mathbf{T} is a lower triangular matrix, the bounds on $\vec{k} = (k_0, k_1, \dots, k_{n-1})$ can be represented in the same manner as the bounds on $\vec{\lambda}^t$ in Eq. (9):

$$\begin{aligned} k_0^{\min} &\leq k_0 \leq k_0^{\max}, \\ k_1^{\min}(k_0) &\leq k_1 \leq k_1^{\max}(k_0), \\ k_2^{\min}(k_0, k_1) &\leq k_2 \leq k_2^{\max}(k_0, k_1), \\ &\dots, \\ k_{n-1}^{\min}(\dots, k_{n-2}) &\leq k_{n-1} \leq k_{n-1}^{\max}(\dots, k_{n-2}). \end{aligned} \quad (13)$$

This representation suggests an algorithm for iterative generation of templates using n nested loops. The outer-most loop generates integer values of k_0 between k_0^{\min} and k_0^{\max} ; for each k_0 , the next inner loop calculates the bounds $k_1^{\min}(k_0)$ and $k_1^{\max}(k_0)$ and generates integer values of k_1 between the two bounds; and so on until the inner-most loop which, for each k_0, \dots, k_{n-2} , generates integer values of k_{n-1} between $k_{n-1}^{\min}(k_0, \dots, k_{n-2})$ and $k_{n-1}^{\max}(k_0, \dots, k_{n-2})$. This algorithm ensures that all of \mathcal{P} is visited, regardless of its geometry, and is the motivation for the representation of \mathcal{P} given by Eq. (9).

We now derive the lower triangular matrix \mathbf{T} which transforms integers \vec{k} to template points $\vec{\lambda}^t$. A n -dimensional lattice embedded in m -dimensional space is




FIG. 4. Histograms of mismatches in μ_{rss} , μ_{ss} , and μ_0 , averaged over all simulated values of T , Δt_{start} , and f_{max} . Only first spindown is used in (a); both first and second spindown are used in (b). An A_n^* lattice is used to generate random parameter offsets with $\mu_{\text{rss}} \leq 0.3$.

generated by a $m \times n$ matrix [33] \mathbf{G} , $m \geq n$, such that

$$\vec{p}^{\ell'} = \mathbf{G}\vec{k}, \quad (14)$$

where $\vec{k} \in \mathbb{Z}^n$ are integer vectors, and $\vec{p}^{\ell'} \in \mathbb{R}^m$ are the lattice points embedded in m -dimensional space. To find a representation of the lattice in n -dimensional space, we compute the QL factorization [34] of \mathbf{G} :

$$\mathbf{G} = \mathbf{Q}\mathbf{L}', \quad (15)$$

where \mathbf{Q} is an $m \times m$ orthogonal matrix, and \mathbf{L}' is an $m \times n$ matrix with zeros above the $(m-n)$ th subdiagonal.

The matrix \mathbf{Q} denotes an overall rotation, which does not affect the lattice's covering properties, and can therefore be discarded. The top $(m-n)$ rows of \mathbf{L}' are zero, and are therefore dropped; let \mathbf{L} denote the lower triangular matrix comprising the remaining, lower n rows of \mathbf{L}' . The lattice can then be generated using

$$\vec{p}^{\ell} = \mathbf{L}\vec{k}, \quad (16)$$

where $\vec{p}^{\ell} \in \mathbb{R}^n$ are the lattice points, now embedded in n -dimensional space.

Finally, we require a transformation from lattice points $\{\vec{p}^{\ell}\}$ to templates $\{\vec{\lambda}^t\}$ such that Eq. (10) is satisfied, i.e. each point $\vec{\lambda} \in \mathcal{P}$ is within a mismatch μ_{max} of some




FIG. 5. Illustration of a two-dimensional parameter space \mathcal{P} (gray shaded area) whose boundaries (solid lines) are represented, per Eq. (9), by $1 \leq \lambda_0^t \leq 2$ and $2/\lambda_0^t \leq \lambda_1^t \leq 5/\lambda_0^t$. Each boundary is labeled by its defining inequality.




FIG. 6. Illustration of the boundary of a two-dimensional parameter space \mathcal{P} . Templates (crosses, dashed ellipses) are laid within the bound $\lambda_1^t \leq \lambda_1^{\text{max}}$ (solid line); the gray shaded area indicates where \mathcal{P} is covered by these templates. The white areas within the $\lambda_1^t \leq \lambda_1^{\text{max}}$ bound are covered by laying templates outside the parameter space (circles, solid ellipses), up to $\lambda_1^t \leq \lambda_1^{\text{max}} + 0.5\beta_1$ (dashed line), where β_1 is the height of the metric ellipse bounding box (solid box).

template $\vec{\lambda}^t$. Let

$$\vec{\lambda} = \frac{\sqrt{\mu_{\text{max}}}}{R} \mathbf{B}\vec{p}, \quad \vec{\lambda}^t = \frac{\sqrt{\mu_{\text{max}}}}{R} \mathbf{B}\vec{p}^{\ell}, \quad (17)$$

where R is the covering radius of the lattice, and \mathbf{B} is the Cholesky factorization of $\mathbf{g}_{\text{rss}}^{-1}$, i.e. a lower triangular matrix satisfying $\mathbf{g}_{\text{rss}}^{-1} = \mathbf{B}\mathbf{B}^T$. Substituting this and

Eq. (17) into Eq. (10) gives

$$\frac{\mu_{\max}}{R^2} [\mathbf{B}(\vec{p} - \vec{p}^\ell)]^T [\mathbf{B}\mathbf{B}^T]^{-1} [\mathbf{B}(\vec{p} - \vec{p}^\ell)] \leq \mu_{\max}, \quad (18)$$

which simplifies to Eq. (6), the definition of a covering sphere (see Section II C). Since Eq. (6) is always satisfied by the lattice points $\{\vec{p}^\ell\}$, Eq. (10) will also always be satisfied by the templates $\{\vec{\lambda}^t\}$. The required lower triangular matrix \mathbf{T} is therefore

$$\mathbf{T} = \frac{\sqrt{\mu_{\max}}}{R} \mathbf{B}\mathbf{L}. \quad (19)$$

Figure 7 plots an example lattice template bank in the reduced supersky metric sky coordinates (n_a, n_b) . The parameter space comprises two unit disks centered on $n_a = \pm 1, n_b = 0$, one for each hemisphere of the sky. To convert from reduced supersky to physical sky coordinates, e.g. right ascension α and declination δ , sky positions are converted to supersky coordinates $\vec{n} = (n_x, n_y, n_z)$, and projected onto the sky sphere $|\vec{n}| = 1$, as illustrated in Fig. 8. Boundary templates which lie outside the reduced supersky parameter space are projected radially onto the parameter-space boundary, as shown in Fig. 7. Finally, (α, δ) are calculated from $\vec{n} = (\cos \alpha \cos \delta, \sin \alpha \cos \delta, \sin \delta)$.

C. Nearest template finding and indexing

Given a lattice template bank $\{\vec{\lambda}^t\}$, we would like to be able to find the nearest template $\vec{\lambda}^t(\vec{\lambda}) \in \{\vec{\lambda}^t\}$ to any given point $\vec{\lambda} \in \mathcal{P}$, from which we can calculate $\vec{\Delta}\lambda = \vec{\lambda} - \vec{\lambda}^t(\vec{\lambda})$ and the mismatch μ_{rss} via Eq. (3). As discussed in Section VI, this facility would also be needed by a semicoherent search. Efficient algorithms for finding $\vec{\lambda}^t(\vec{\lambda})$ are specialized to the type of lattice being used. For example, the most efficient algorithm for \mathbb{Z}^n lattices is

$$\vec{\lambda}^t(\vec{\lambda})|_{\mathbb{Z}^n} = \mathbf{T} \left\lfloor \mathbf{T}^{-1} \vec{\lambda} \right\rfloor, \quad (20)$$

where the operation $\lfloor \cdot \rfloor$ rounds each vector element to the nearest integer. For A_n^* lattices, several efficient algorithms have been proposed; see [22, 35] and references therein.

We implement the algorithm described [36] in [35], as it gives the best known scaling (linear) with the lattice dimension n . Essentially, the algorithm maps points in the space of A_n^* to points in the space of \mathbb{Z}^{n+1} , finds the nearest point in \mathbb{Z}^{n+1} using Eq. (20), then efficiently determines which point in A_n^* this point corresponds to. The algorithm takes as input an $(n+1)$ -dimensional vector $(0, \mathbf{T}^{-1} \vec{\lambda})$ and returns an $(n+1)$ -dimensional vector $\vec{k}' \in \mathbb{Z}^{n+1}$, from which the nearest template is given by

$$\vec{k} = (k'_1 - k'_0, k'_2 - k'_0, \dots, k'_n - k'_0), \quad (21)$$

$$\vec{\lambda}^t(\vec{\lambda})|_{A_n^*} = \mathbf{T} \vec{k}. \quad (22)$$

We also would like to have an efficient lookup table from every template $\vec{\lambda}^t \in \{\vec{\lambda}^t\}$ to a unique index $j(\vec{\lambda}^t) = 0, 1, \dots, \mathcal{N} - 1$, where \mathcal{N} is the number of templates in $\{\vec{\lambda}^t\}$. For example, $j(\vec{\lambda}^t)$ might index an array of the \mathcal{N} values of the \mathcal{F} -statistic computed at each template. The lookup table would also be needed when implementing a semicoherent search, as discussed in Section VI.

We implement the $j(\vec{\lambda}^t)$ lookup table using a *trie*, a tree data structure also known as a digital tree [e.g. 37]. A trie has the advantage, compared to a hash table, of a constant lookup time of order n , and no possibility of key collisions which may degrade performance. Its operation is illustrated in Fig. 9. The trie takes the vector \vec{k} given by Eq. (21) as input. First, k_0 determines which edge to follow from the input node; next, k_1 determines which edge to follow from the chosen k_0 node; k_2 then determines which edge to follow from the chosen k_1 node; and so on up to k_{n-2} . From the chosen k_{n-2} node, values for the k_{n-1} lower bound k_{n-1}^{\min} and its index j_{n-1}^{\min} are retrieved. Together with k_{n-1} , the index $j(\vec{\lambda}^t)$ is then given by

$$j(\vec{\lambda}^t) = j_{n-1}^{\min} + (k_{n-1} - k_{n-1}^{\min}). \quad (23)$$

For example, for the example trie in Fig. 9, a template with $\vec{k} = (-42, 5, 23, \dots, 11)$ would have $k_{n-1}^{\min} = 9$, $j_{n-1}^{\min} = 12$, and hence $j(\vec{\lambda}^t) = 12 + (11 - 9) = 14$. The lookup trie is constructed by generating each template in the bank, and filling the trie with the k_0, \dots, k_{n-2} of each template, and the k_{n-1}^{\min} and j_{n-1}^{\min} of templates in the $(n-1)$ th dimension.

D. Lattice template placement testing

This section presents tests performed on lattice template banks generated by the algorithm described in the previous sections. Details of input parameters to the tests are summarized in Table I.

Template banks are generated using four types of parameter spaces, labeled A_1 , A_2 , B_1 , and B_2 . The maximum mismatches μ_{\max} used for each type, and the list of time spans T for which templates banks of each type are generated, are given in Table I. All types are also generated at five values of start time $\Delta t_{\text{start}} = t_0 \pm \{0, 90, 180\}$ days, where t_0 is given in Section III B.

Parameter spaces A_X ($X = 1, 2$) cover the whole sky, a fixed band in reduced supersky frequency ν of width 10^{-6} Hz, and a fixed band in reduced supersky spindown $\dot{\nu}$ of $[-10^{-9}, 0]$ Hz s $^{-1}$. The B_X parameter spaces cover the whole sky, a fixed band in physical frequency f of width 10^{-6} Hz, and a fixed band in physical spindown \dot{f} of $[-10^{-9}, 0]$ Hz s $^{-1}$. The frequency bands start at 100 Hz for A_1 and B_1 , and at 500 Hz for A_2 and B_2 . The bandwidths are limited by the computational cost of generating the template banks and calculating the mismatches to each test point.




FIG. 7. Example of a lattice template bank in reduced supersky coordinates (n_a, n_b) , with $T = 1$ day, $\Delta t_{\text{start}} = 0$, and $\nu_{\text{max}} = 100$ Hz. Templates and their metric ellipses are plotted as points and dashed lines respectively. Templates, plotted as crosses, which lie outside the parameter-space boundaries (circles) are moved radially onto the nearest boundary (thick lines) when converted to physical coordinates (α, δ) . An A_n^* lattice is used to generate the template bank with $\mu_{\text{rss}} \leq 0.3$.

TABLE I. Input parameters to tests of the lattice template placement algorithm (Section IV); properties of the template banks generated by the tests, averaged over Δt_{start} at fixed T ; and errors in estimates of the number of templates (Section V).

Quantity	Value for parameter space														
	A ₁				A ₂				B ₁			B ₂			
Starting frequency / Hz	$\nu = 100$				$\nu = 500$				$f = 100$			$f = 500$			
Frequency band / Hz	$\Delta\nu = 10^{-6}$				$\Delta\nu = 10^{-6}$				$\Delta f = 10^{-6}$			$\Delta f = 10^{-6}$			
First spindown band / Hz	$\Delta\dot{\nu} = 10^{-9}$				$\Delta\dot{\nu} = 10^{-9}$				$\Delta\dot{f} = 10^{-9}$			$\Delta\dot{f} = 10^{-9}$			
Maximum mismatch	$\mu_{\text{max}} = 0.3$				$\mu_{\text{max}} = 0.6$				$\mu_{\text{max}} = 0.3$			$\mu_{\text{max}} = 0.6$			
Time span T / days	1	3	9	27	1	3	9	27	1	3	9	1	3	9	
$\log_{10}(\text{Number of templates } \mathcal{N})$	6.3	6.8	7.2	8.7	7.3	7.7	8.1	9.6	7.8	8.2	9.4	8.8	9.1	10.4	
$\log_{10}(\text{Bulk templates}/\mathcal{N})$	−3	−2	−1	−0.8	−3	−2	−1	−0.9	−5	−4	−4	−5	−4	−4	
$\log_{10}(\text{Missed test points}/\mathcal{N})$	−∞	−∞	−∞	−∞	−∞	−∞	−∞	−∞	−∞	−16	−17	−∞	−14	−15	
Mean reduced supersky $\langle\mu_{\text{rss}}\rangle$	0.16	0.16	0.16	0.16	0.31	0.31	0.31	0.31	0.16	0.16	0.16	0.31	0.31	0.31	
Mean supersky $\langle\mu_{\text{ss}}\rangle$	0.15	0.15	0.16	0.16	0.31	0.31	0.31	0.32	0.15	0.15	0.16	0.31	0.31	0.31	
$\langle \mathcal{N}_{\text{est}}(\text{num. } \mathbf{g}_{\text{rss}}) - \mathcal{N} /\mathcal{N}\rangle / \%$	1.8	1.3	2.6	5.6	4.3	2.4	3.8	2.3	8.1	1.5	1.8	3.4	0.3	0.5	
$\max \mathcal{N}_{\text{est}}(\text{num. } \mathbf{g}_{\text{rss}}) - \mathcal{N} /\mathcal{N} / \%$	2.3	4.8	4.7	7.5	6.8	3.6	16.5	2.9	8.4	1.5	2.1	3.5	0.3	0.6	
$\langle \mathcal{N}_{\text{est}}(\text{appx. } \mathbf{g}_{\text{rss}}) - \mathcal{N} /\mathcal{N}\rangle / \%$	1.7	4.1	13.8	6.0	3.4	4.9	17.8	3.8	78.9	44.7	48.8	80.5	45.5	49.1	
$\max \mathcal{N}_{\text{est}}(\text{appx. } \mathbf{g}_{\text{rss}}) - \mathcal{N} /\mathcal{N} / \%$	2.2	7.7	17.0	10.0	6.0	6.1	33.7	6.1	88.5	68.8	73.4	89.4	69.2	73.8	

For each template $\tilde{\lambda}^t$ in each generated template bank, ~ 10 test points $\tilde{\lambda}$ are randomly drawn from the parameter space, and the nearest template $\tilde{\lambda}^t(\tilde{\lambda})$ to each test point is found as described in Section IV C. The mismatches μ_{rss} and μ_{ss} between each test point and their

nearest templates are computed via Eq. (3) for both reduced supersky and supersky metrics.

Table I lists, for each parameter space and each T , the following properties computed by the tests: the number of templates; the fraction of “bulk” templates, i.e. ex-




FIG. 8. The projection of the example lattice template bank from Fig. 7 onto the sky sphere $|\vec{n}| = 1$. Templates are plotted as points; lines are drawn from each template to its (up to) six nearest neighbors in the A_n^* lattice.

cluding extra templates needed to cover the boundaries (see Fig. 6 and Section IV E); the fraction of missed test points, i.e. where $\mu_{\text{rss}} > \mu_{\text{max}}$; and the means of the mismatch distributions of μ_{rss} and μ_{ss} . These quantities are averaged over the five start times Δt_{start} .

Due to the limited frequency and spindown bandwidths, the template banks are dominated by boundary templates, as evidenced by the relatively small fractions of bulk templates (Table I). Since the simulations in Section III effectively test the properties of an infinite template bank, i.e. without boundaries, it is complementary that the tests presented here test the properties of template banks dominated by boundaries.

The fractions of missed test points (Table I) indicate “holes” in the template bank, i.e. regions of parameter space not covered by templates to within the desired maximum mismatch. No holes were found in parameter spaces A_X , i.e. no test points were missed, and only very small holes were found in B_X , the number of missed test points per template being $\lesssim 10^{-14}$.

The mean mismatches $\langle \mu_{\text{rss}} \rangle$ and $\langle \mu_{\text{ss}} \rangle$, shown in Table I, are similar for the reduced supersky and supersky metrics respectively. Figure 10 plots the averaged mismatch distributions of μ_{rss} and μ_{ss} for each parameter space, which also show close agreement. While the \mathcal{F} -statistic mismatch was not computed in these tests, the similarity of the histograms in Fig. 10 to those in Fig. 4a leads us to expect (were it to be computed) a \mathcal{F} -statistic mismatch histogram similar to that in Fig. 4a.




FIG. 9. Illustration of the $j(\vec{\lambda}^t)$ lookup trie described in Section IV C. The trie is traversed top to bottom. Ovals denote the input and output nodes of the trie. Diamonds denote nodes where a branch is taken. Rectangles denote nodes where a stored value is retrieved. Cloud shapes denote nodes omitted for brevity.

E. The “staircase” boundary template issue

This section describes an issue that is encountered when covering the boundaries of certain parameter spaces, in particular the parameter spaces B_X of the previous section. We call it the “staircase” issue. It arises because the boundaries of the parameter space, which are represented per Eq. (9) by continuous functions $\lambda^{\min}(\dots)$ and $\lambda^{\max}(\dots)$, are being covered by a bank of discrete templates $\{\lambda^t\}$. If, in parameter-space dimension i , the $\lambda_i^{\min}(\dots)$ and $\lambda_i^{\max}(\dots)$ functions change significantly on the scale of a single template, e.g. the width of its metric ellipse bounding box β_i , regions of the parameter space boundary in dimension i may not be covered.

Figure 11 illustrates an example of the staircase issue. The example parameter space of Fig. 5, shaded light gray, is covered with a template bank which is deliberately coarse in the λ_0^t coordinate, comprising only four columns of templates; the boundary templates of these columns are plotted as crosses. For simplicity, we take the region covered by this template bank to be the union of the bounding boxes around each template; this gives




FIG. 10. Histograms of mismatches in μ_{rss} and μ_{ss} , generated by the lattice template placement algorithm detailed in Section IV, and averaged over all values of T and Δt_{start} . The plots correspond to the parameter spaces (a) A_1 , (b) A_2 , (c) B_1 , and (d) B_2 from Table I. An A_n^* lattice is used to place templates with $\mu_{\text{rss}} \leq 0.3$ (left column) and $\mu_{\text{rss}} \leq 0.6$ (right column).




FIG. 11. Illustration of the staircase boundary template issue, using the example parameter space of Fig. 5. The parameter space to be covered is light gray. The short-dashed outline shows the area covered by the template bank, if the bounds $\lambda_1^{\min}(\lambda_0^t)$ and $\lambda_1^{\max}(\lambda_0^t)$ are computed exactly at the boundary templates (crosses). Regions of the parameter space not covered by this template bank are dark gray. The long-dashed outline shows the area covered if instead the bounds are extremized over the bounding box of each boundary template.

the short-dashed, staircase-shaped outline. (For the purpose of this illustration, we ignore the extra boundary templates illustrated in Fig. 6.)

It is clear that there are regions of the parameter space, shaded dark gray, not covered by this template bank. It is also clear that the uncovered regions exist because the short-dashed outline is a poor approximation to the true parameter-space shape. A much denser template bank in λ_0^t , so that the bound functions $\lambda_1^{\min}(\lambda_0^t)$ and $\lambda_1^{\max}(\lambda_0^t)$ are better approximated, would not suffer from this issue. By analogy, imagine trying to approximate the area under a curve $f(x)$ by a series of rectangles of fixed width δx and centered at points x_j , i.e. $\int f(x)dx \approx \sum_j f(x_j)\delta x$. We would expect a poor approximation if $f(x_j + \delta x) \not\approx f(x_j)$ for many x_j . Conversely, one would expect a good approximation if $f(x_j + \delta x) \approx f(x_j)$ for most x_j , which may be achieved simply by reducing δx and adding more points x_j .

In the case of template placement, however, it would be undesirable to increase the template bank density, since this would inflate the computational cost of the search. We can, instead, improve how well the region covered by the template bank approximates the true shape of the parameter space. Instead of calculating the bound functions $\lambda_i^{\min}(\dots)$ and $\lambda_i^{\max}(\dots)$ solely at a boundary template, we can find their extrema over the region covered by that template, e.g. its bounding box. This ensures that the template bank extends far enough to cover any

variations in the boundaries that might occur between neighboring boundary templates.

For example, in Fig. 11, we can replace the functions $\lambda_1^{\min}(\lambda_0^t)$ and $\lambda_1^{\max}(\lambda_0^t)$ with the following:

$$\lambda_1^{\min}(\lambda_0) \rightarrow \min_{|\delta\lambda_0| \leq 0.5\beta_0} \lambda_1^{\min}(\lambda_0 + \delta\lambda_0), \quad (24a)$$

$$\lambda_1^{\max}(\lambda_0) \rightarrow \max_{|\delta\lambda_0| \leq 0.5\beta_0} \lambda_1^{\max}(\lambda_0 + \delta\lambda_0), \quad (24b)$$

where β_0 is the bounding box width in the λ_0 coordinate. Doing so gives the long-dashed outline in Fig. 11, which now completely covers the parameter space.

The generalization of Eqs. (24) to any parameter-space bound function is, e.g. for $\lambda_i^{\min}(\dots)$:

$$\lambda_i^{\min}(\lambda_0, \dots, \lambda_{i-1}) \rightarrow \min_{\substack{|\delta\lambda_0| \leq 0.5\beta_0 \\ \vdots \\ |\delta\lambda_{i-1}| \leq 0.5\beta_{i-1}}} \lambda_i^{\min}(\lambda_0 + \delta\lambda_0, \dots, \lambda_{i-1} + \delta\lambda_{i-1}). \quad (25)$$

If it is known that the bound functions are monotonic over the extent of the bounding box, the extrema can only occur at the vertices of the bounding box. It is then sufficient to extremize the bounds only over the finite set of vertices, e.g.:

$$\lambda_i^{\min}(\lambda_0, \dots, \lambda_{i-1}) \rightarrow \min \lambda_i^{\min}(\lambda_0 \pm 0.5\beta_0, \dots, \lambda_{i-1} \pm 0.5\beta_{i-1}). \quad (26)$$

The staircase issue appears when covering the parameter spaces B_X of Section IV D. While these parameter spaces cover fixed bands in physical frequency/spindown $f^{(s)}$, the parameter-space bounds are specified as functions of the reduced supersky frequencies/spindowns $\nu^{(s)}$. The two coordinate systems are related by (see Paper I)

$$\nu^{(s)} = f^{(s)} + \vec{\Delta}^s \cdot \vec{n}, \quad (27)$$

where the $\vec{\Delta}^s$ are offset vectors found when computing the reduced supersky metric. When covering a fixed band in $f^{(s)}$, the bounds on $\nu^{(s)}$, e.g. $\nu^{(s)\min}(\vec{n})$, are therefore functions of sky position \vec{n} .

The difference in $\nu^{(s)\min}(\vec{n})$ between neighboring sky positions \vec{n}_1 and \vec{n}_2 is

$$|\nu^{(s)\min}(\vec{n}_1) - \nu^{(s)\min}(\vec{n}_2)| = |\vec{\Delta}^s \cdot (\vec{n}_1 - \vec{n}_2)| \sim |\Delta_a^s| \beta_a + |\Delta_b^s| \beta_b, \quad (28)$$

where $\vec{\beta}$ is the metric ellipse bounding box given by Eq. (11). Depending on the density of sky templates, this difference can be much larger than the extent of the bounding box in $\nu^{(s)}$. For example, when $T \sim 1$ day, $|\nu^{\min}(\vec{n}_1) - \nu^{\min}(\vec{n}_2)| \sim 10^{-3}$ Hz, whereas $\beta_\nu \sim 10^{-5}$ Hz. In short, the bounds on ν are changing on a much larger scale ($\sim 10^{-3}$ Hz) than the extent of a single template ($\sim 10^{-5}$ Hz); this is precisely the conditions where the staircase issue becomes important.

At the time the tests in Section IV D were performed, the solution to the staircase issue outlined above [i.e. Eq. (25)] was not realized. Instead, an empirical solution was used, where the bounds on $\nu^{(s)}$ were simply extended by $|\Delta_a^s| \beta_a + |\Delta_b^s| \beta_b$ [Eq. (28)]. This solution results in only a small fraction of missed test points in the B_X parameter spaces (Table I). The solution suggested by Eq. (26) will however be used in future implementations.

V. NUMBER OF TEMPLATES

Having confirmed, in the previous section, that lattice template placement using the reduced supersky metric can be successfully implemented, we now investigate the number of templates required to cover an all-sky parameter space. Section V A considers how to accurately estimate the number of templates; this is needed in order to determine the most sensitive semicoherent search setup [e.g. 20]. Section V B examines the scaling of the number of templates with time span T , which determines the relationship between search sensitivity and computational cost.

A. Counted and estimated number of templates

The number of templates required to cover an n -dimensional parameter space \mathcal{P} with constant metric \mathbf{g} , using a lattice template bank with maximum mismatch μ_{\max} , is estimated by [23, 38]

$$\mathcal{N}_{\text{est}} = \theta \mu_{\max}^{-n/2} \sqrt{\det \mathbf{g}} \int_{\mathcal{P} \cup \partial \mathcal{P}} d\vec{\lambda}, \quad (29)$$

where θ is the normalized thickness of the lattice being used (see Section II C). The integral gives the volume, with respect to the parameters $\vec{\lambda}$, of both \mathcal{P} and its boundary $\partial \mathcal{P}$, to account for the extra boundary templates discussed in Sections IV A and IV E.

The determinant of the reduced supersky metric may be written as

$$\det \mathbf{g}_{\text{rss}} = \det \mathbf{g}_{\text{rss}, \vec{n} \vec{n}} \cdot \det \mathbf{g}_{\text{rss}, \nu^{(s)} \nu^{(s)}}, \quad (30)$$

where $\mathbf{g}_{\text{rss}, \vec{n} \vec{n}}$ is the sky-sky block and $\mathbf{g}_{\text{rss}, \nu^{(s)} \nu^{(s)}}$ the frequency-frequency block of \mathbf{g}_{rss} , the off-diagonal sky-frequency blocks being zero. The matrix $\mathbf{g}_{\text{rss}, \nu^{(s)} \nu^{(s)}}$ is identical to the physical frequency/spindown metric, for which analytic expressions exist [e.g 39–41]. Its determinant is therefore known exactly, e.g.:

$$\det \mathbf{g}_{\text{rss}, \nu^{(s)} \nu^{(s)}} = \begin{cases} \frac{\pi^4 T^6}{540}, & s_{\max} = 1, \\ \frac{\pi^6 T^{12}}{13608000}, & s_{\max} = 2, \end{cases} \quad (31)$$

where s_{\max} is the maximum number of spindowns. The matrix $\mathbf{g}_{\text{rss}, \vec{n} \vec{n}}$ is diagonal, and its determinant is therefore the product of its diagonal elements $g_{n_a n_a}$ and $g_{n_b n_b}$.

In turn, these are reasonably well-approximated by functions of T only, ignoring the weaker dependence on t_{start} due to the Earth's noncircular orbital motion. Analytic expressions approximating $g_{n_a n_a}$ and $g_{n_b n_b}$ are given by Eqs. (A1) in the Appendix.

Table I shows errors $|\mathcal{N}_{\text{est}} - \mathcal{N}|/\mathcal{N}$ between the number of templates \mathcal{N} counted in the template banks generated in Section IV, and the number \mathcal{N}_{est} estimated using Eq. (29). The \mathbf{g}_{rss} used in Eq. (29) is either computed numerically, or approximated using Eqs. (30), (31), and (A1). Using the numerically-computed \mathbf{g}_{rss} , errors are a few percent on average, and are limited to $\lesssim 17\%$ over the four types of parameter spaces. This confirms that the behavior of the algorithm of Section IV conforms to that expected by Eq. (29). The approximate \mathbf{g}_{rss} also leads to reasonable errors of $\lesssim 34\%$, except for parameter spaces B_X , where they are between 44% and 90%. This is due to using typical values for the elements of the offset vectors $\vec{\Delta}^s$ when estimating the number of extra staircase boundary templates: $|\Delta_a^0| \sim |\Delta_b^0| \sim 5 \times 10^{-5} f_{\text{max}}$ and $|\Delta_a^1| \sim |\Delta_b^1| \sim 10^{-11} f_{\text{max}}$. Nevertheless, we imagine that the approximate method may still be useful for rapid order-of-magnitude template counting, e.g. during search setup optimization [e.g. 20].

Figure 12 plots the number of templates, as functions of T , required to search a more realistic parameter space than those used for testing in Section IV: the whole sky, a frequency band of 10^{-2} Hz at $f = 100$ Hz, and bands of widths 10^{-8} Hz s^{-1} in first spindown and 10^{-18} Hz s^{-2} in second spindown. (In particular, the frequency band is more typical of an Einstein@Home “work unit” [8]) The number of templates is counted by the algorithm of Section IV, and estimated using Eq. (29). Two estimates are performed: using the numerically-computed \mathbf{g}_{rss} and taking into account the extra staircase boundary templates discussed in Section IV E, and using the approximate \mathbf{g}_{rss} and ignoring the extra staircase boundary templates. Good agreement is seen between the number of counted templates and the various estimates. For first spindown-only searches, the fraction of templates required to address the staircase issue steadily decreases with T ; when second spindown is added, the fraction is larger, indicating that the second spindown band is more dominated by its boundaries.

Brady *et al.* [15] predict the number of templates required to search the whole sky, frequencies up to 200 Hz, and spindowns $|\dot{f}| \leq f/10^3$ yr, motivated by a minimum “spindown age” of the gravitational-wave pulsar. Their estimate, given by their Eqs. (6.3)–(6.7) and (6.9), used a different derivation of the parameter-space metric (based on loss of power, not loss of \mathcal{F} -statistic) which is nonconstant; it is therefore an idealized estimate of the minimal number of templates achievable. It assumed an hexagonal prism lattice (i.e. the composite lattice $A_2^* \otimes \mathbb{Z}^{n-2}$), and a maximum *projected* mismatch of 0.3 (i.e. assuming mismatch has already been minimized over frequency).

Figure 13 compares the estimate of [15] to an estimate given by Eq. (29), using the approximate $\det \mathbf{g}_{\text{rss}}$, an A_n^*




FIG. 12. Number of templates required to search a coherent segment of time span T , centered on the reference time t_0 of Section III B, over the following parameter spaces, using an A_n^* lattice template bank with $\mu_{\text{max}} = 0.3$: (a) the whole sky, a fixed band in f of $[100, 100.01]$ Hz, and a fixed band in \dot{f} of $[-10^{-8}, 0]$ Hz s^{-1} ; (b) as above, with an additional fixed band in \ddot{f} of $[0, 10^{-18}]$ Hz s^{-2} . Crosses denote the number \mathcal{N} counted by the algorithm of Section IV; lines denote the number \mathcal{N}_{est} estimated by the procedures described in Section V. Black/gray denotes the template counts including/excluding the extra boundary templates discussed in Section IV E.

lattice, and a maximum (total) mismatch of $\mu_{\text{max}} = 0.3$. The two estimates are in close agreement, differing by less than a factor of two. Given the idealized assumptions of [15], it is encouraging to see close agreement with the template estimation described in this section, which in turn is in good agreement with the practical implementation of lattice template placement described in Section IV.




FIG. 13. Number of templates required to search a coherent segment of time span T over the whole sky, a fixed band in f of $[0, 200]$ Hz, and frequency-dependent bands in \dot{f} of $\pm f/(10^3 \text{ yr})$. The estimate of Brady *et al.* [15] (dashed line) is compared with an estimate following the procedure of Section V using the approximate \mathbf{g}_{rss} (solid line).

B. Scaling of number of templates with time span

The number of templates \mathcal{N} may be modeled, in the vicinity of a fixed time span T_0 , by a power law in T : $\mathcal{N}(T)|_{T=T_0} \approx T^p$, where the local power-law exponent is

$$p = \left. \frac{d[\log \mathcal{N}(T)]}{d[\log T]} \right|_{T=T_0}. \quad (32)$$

This exponent is used to relate the sensitivity of a search and its computational cost [20]. The scaling of the number of frequency/spindown templates $\mathcal{N}_{\text{freq}}$ is known exactly: the number of templates in each frequency/spindown coordinate $\nu^{(s)}$ scales as T^{s+1} , and hence $\mathcal{N}_{\text{freq}} \propto T^{(s_{\text{max}}+1)(s_{\text{max}}+2)/2}$. The scaling of the number of sky templates \mathcal{N}_{sky} is, however, less straightforward. It is generally stated that $\mathcal{N}_{\text{sky}} \propto T^q$ with $q \geq 2$ [1, 15, 21, 39, 42, 43]; on the other hand, [18] reported \mathcal{N}_{sky} to be approximately constant once $T \gtrsim 2$ days.

Figure 14 plots the local power-law exponent p of the number of templates $\mathcal{N} = \mathcal{N}_{\text{sky}}\mathcal{N}_{\text{freq}}$. When including only first spindown, we find a similar scaling of \mathcal{N}_{sky} to that of [18]; $p \approx 3$ for $2 \lesssim T \lesssim 10$ days, which implies (since $\mathcal{N}_{\text{freq}}$ is known to scale as T^3) that \mathcal{N}_{sky} is approximately constant over this period. Once $T \gtrsim 10$ days, however, \mathcal{N}_{sky} scales with an increasing power $q \lesssim 3$. When including both first and second spindown, the scaling of \mathcal{N} with T becomes more complicated, which is likely due to the Earth’s noncircular orbital motion (see Paper I).




FIG. 14. Local power-law exponent p of the number of templates plotted in Figs. 12a, with first spindown (solid line), and 12b (dashed line), with both first and second spindown. The thick horizontal lines denote the expected p , assuming that $\mathcal{N}_{\text{sky}} \propto T^2$ and $\mathcal{N}_{\text{freq}} \propto T^3$ (solid) or T^6 (dashed).

VI. DISCUSSION

This paper demonstrates, for the first time, practical lattice template placement for an all-sky broadband-frequency search for gravitational-wave pulsars. The mismatch predictions of the reduced supersky metric, derived previously in Paper I, are confirmed using realistic mismatch distributions, as would be encountered in a real search. An algorithm implementing lattice template placement, including efficient iteration over the parameter space, nearest template finding, and correct treatment of the parameter-space boundaries, is described and tested. The number of templates generated by the algorithm is consistent with theoretical expectations, and with previous results in the literature.

Future work will extend the reduced supersky metric to a semicoherent search. In such a search, the data are partitioned into N data segments, and an average \mathcal{F} -statistic is computed on a high-resolution “fine” template bank from N \mathcal{F} -statistic values, which are computed by coherently match-filtering each segment on N low-resolution “coarse” template banks. Typically, the template in each coarse template bank which is closest to the current fine template is chosen to contribute to the average \mathcal{F} -statistic; this requires an algorithm for finding the nearest template, as that described in Section IV C. The resolution of the fine template bank is determined by the parameter-space metric of the average \mathcal{F} -statistic, which is the average of the parameter-space metrics of each data segment [44]. The extension of the reduced supersky metric to an averaged metric will need to be investigated.

While the focus of this paper has been searches for isolated gravitational-wave pulsars, the lattice template al-

gorithm presented here may be applied to any parameter-space described by a constant metric. An interesting example is searches for gravitational-wave pulsars in low-mass X-ray binary systems with a known sky location [45], where the parameter space comprises frequency and the orbital parameters of the binary system. Recent work on the parameter-space metric [46] focuses on the feasibility of a search for Scorpius X-1, one of the most promising gravitational-wave sources of this type.

ACKNOWLEDGMENTS

I thank Reinhard Prix for many valuable discussions. Numerical simulations were performed on the ATLAS computer cluster of the Max-Planck-Institut für Gravitationsphysik. This paper has document numbers AEI-2014-052 and LIGO-P1400202.

Appendix A: Approximations to the reduced supersky metric sky elements

The elements $g_{n_a n_a}$ and $g_{n_b n_b}$ of the reduced supersky metric \mathbf{g}_{rss} are approximated by the following expressions, where $\log T \equiv \log(T/1 \text{ s})$: for first spindown,

$$\log \frac{g_{n_a n_a}}{f_{\text{max}}^2} \approx \begin{cases} \frac{(\log T)^2}{3.803} - \frac{\log T}{0.2836}, & 0.5 \leq \frac{T}{\text{day}} < 1.38, \\ \frac{(\log T)^2}{19.04} - \frac{\log T}{0.9374}, & 1.38 \leq \frac{T}{\text{day}} < 10.5, \\ \frac{(\log T)^2}{2.669} - \frac{\log T}{0.1823}, & 10.5 \leq \frac{T}{\text{day}} < 60, \end{cases} \quad \log \frac{g_{n_b n_b}}{f_{\text{max}}^2} \approx \begin{cases} \frac{(\log T)^2}{7.751} - \frac{\log T}{0.4736}, & 0.5 \leq \frac{T}{\text{day}} < 0.875, \\ \frac{(\log T)^2}{3.245} - \frac{\log T}{0.2425}, & 0.875 \leq \frac{T}{\text{day}} < 1.75, \\ \frac{(\log T)^2}{24.81} - \frac{\log T}{1.076}, & 1.75 \leq \frac{T}{\text{day}} < 29, \\ \frac{(\log T)^2}{2.5} - \frac{\log T}{0.1605}, & 29 \leq \frac{T}{\text{day}} < 60; \end{cases} \quad (\text{A1a})$$

and for second spindown,

$$\log \frac{g_{n_a n_a}}{f_{\text{max}}^2} \approx \begin{cases} \frac{(\log T)^2}{7.629} - \frac{\log T}{0.4687}, & 0.5 \leq \frac{T}{\text{day}} < 0.875, \\ \frac{(\log T)^2}{3.45} - \frac{\log T}{0.2553}, & 0.875 \leq \frac{T}{\text{day}} < 1.88, \\ \frac{(\log T)^2}{26.62} - \frac{\log T}{1.123}, & 1.88 \leq \frac{T}{\text{day}} < 28, \\ \frac{(\log T)^2}{2.591} - \frac{\log T}{0.1663}, & 28 \leq \frac{T}{\text{day}} < 60, \end{cases} \quad \log \frac{g_{n_b n_b}}{f_{\text{max}}^2} \approx \begin{cases} \frac{(\log T)^2}{12.64} - \frac{\log T}{0.6331}, & 0.5 \leq \frac{T}{\text{day}} < 1, \\ \frac{(\log T)^2}{3.194} - \frac{\log T}{0.2359}, & 1 \leq \frac{T}{\text{day}} < 2.12, \\ \frac{(\log T)^2}{22.67} - \frac{\log T}{1.021}, & 2.12 \leq \frac{T}{\text{day}} < 60. \end{cases} \quad (\text{A1b})$$

These are derived by numerically computing \mathbf{g}_{rss} at a fixed f_{max} , and over the following ranges of time span T and start time t_{start} : T from 0.5 to 3 days in steps of 0.125 days, then from 3 to 60 days in steps of 0.5 days; $t_{\text{start}} = t_0 + \Delta t_0 - 0.5T$, where t_0 is one of UTC 2000-06-18 23:59:47, 2009-01-05 11:59:45, or 2017-07-24 23:59:44, and Δt_0 from 0 to 720 days in steps of 1 day. Sinusoidal modulations in $g_{n_a n_a}$ and $g_{n_b n_b}$ with respect to t_{start} , with periods of a synodic month and a year, were $\lesssim 10\%$. The two expressions $\log(\langle g_{\sigma\sigma} \rangle / f_{\text{max}}^2) / \log T$, where $\sigma \in \{n_a, n_b\}$ and $\langle \cdot \rangle$ denotes averaging over t_{start} , are then fitted by piecewise linear functions in $\log T$.

-
- [1] R. Prix, in *Neutron Stars and Pulsars*, Astrophysics and Space Science Library, Vol. 357, edited by W. Becker (Springer, Berlin/Heidelberg, 2009) p. 651, preprint available from <https://dcc.ligo.org/LIGO-P060039/public>.
 - [2] B. Sathyaprakash and B. F. Schutz, Living Reviews in Relativity **12** (2009), 10.12942/lrr-2009-2, retrieved 2014-10-18.
 - [3] B. J. Owen, Physical Review Letters **95**, 211101 (2005), arXiv:astro-ph/0503399.
 - [4] N. K. Johnson-McDaniel and B. J. Owen, Physical Review D **88**, 044004 (2013), arXiv:1208.5227 [astro-ph.SR].
 - [5] B. P. Abbott *et al.* (LIGO Scientific Collaboration), Reports on Progress in Physics **72**, 076901 (2009), arXiv:0711.3041 [gr-qc].
 - [6] T. Accadia *et al.* (Virgo Collaboration), Journal of Instrumentation **7**, P03012 (2012).
 - [7] J. Abadie *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Physical Review D **85**, 022001 (2012), arXiv:1110.0208 [gr-qc].
 - [8] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Physical Review D **87**, 042001 (2013).
 - [9] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), The Astrophysical Journal **785**, 119 (2014).
 - [10] G. M. Harry *et al.* (LIGO Scientific Collaboration), Classical and Quantum Gravity **27**, 084006 (2010).
 - [11] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), arXiv (2013), 1304.0670.
 - [12] K. Somiya (KAGRA Collaboration), Classical and Quantum Gravity **29**, 124007 (2012), arXiv:1111.7185 [gr-qc].
 - [13] P. Jaranowski, A. Królak, and B. F. Schutz, Physical Review D **58**, 063001 (1998), arXiv:gr-qc/9804014.
 - [14] C. Cutler and B. F. Schutz, Physical Review D **72**, 063006 (2005), arXiv:gr-qc/0504011.
 - [15] P. R. Brady, T. Creighton, C. Cutler, and B. F. Schutz, Physical Review D **57**, 2101 (1998), arXiv:gr-qc/9702050.

- [16] B. Krishnan, A. M. Sintes, M. A. Papa, B. F. Schutz, S. Frasca, and C. Palomba, *Physical Review D* **70**, 082001 (2004), arXiv:gr-qc/0407001.
- [17] V. Dergachev, *Classical and Quantum Gravity* **27**, 205017 (2010).
- [18] H. J. Pletsch, *Physical Review D* **82**, 042002 (2010), arXiv:1005.0395 [gr-qc].
- [19] C. Cutler, I. Gholami, and B. Krishnan, *Physical Review D* **72**, 042004 (2005), arXiv:gr-qc/0505082.
- [20] R. Prix and M. Shaltev, *Physical Review D* **85**, 084010 (2012), arXiv:1201.4321 [gr-qc].
- [21] R. Prix, *Physical Review D* **75**, 023004 (2007), arXiv:gr-qc/0606088.
- [22] J. H. Conway and N. J. A. Sloane, *Sphere Packings, Lattices and Groups*, Grundlehren der mathematischen Wissenschaften No. 290 (Springer-Verlag, New York, 1988).
- [23] R. Prix, *Classical and Quantum Gravity* **24**, S481 (2007), arXiv:0707.0428 [gr-qc].
- [24] K. Wette and R. Prix, *Physical Review D* **88**, 123005 (2013), arXiv:1310.5587 [gr-qc].
- [25] P. Jaranowski and A. Królak, *Physical Review D* **59**, 063003 (1999), arXiv:gr-qc/9809046.
- [26] P. Astone, K. M. Borkowski, P. Jaranowski, and A. Królak, *Physical Review D* **65**, 042003 (2002), arXiv:gr-qc/0012108.
- [27] H. J. Pletsch and B. Allen, *Physical Review Letters* **103**, 181102 (2009), arXiv:0906.0023 [gr-qc].
- [28] C. Messenger, R. Prix, and M. A. Papa, *Physical Review D* **79**, 104017 (2009), arXiv:0809.5223 [gr-qc].
- [29] N. J. Higham, *Accuracy and Stability of Numerical Algorithms*, 2nd ed. (SIAM, Philadelphia, 2002).
- [30] The implementation is available as the LATTICETILING module of the LALPULSAR library, a part of the LALSUITE software for gravitational-wave data analysis; see <https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html>.
- [31] K. W. Wette, *Gravitational waves from accreting neutron stars and Cassiopeia A*, Ph.D. thesis, The Australian National University (2009).
- [32] J. Abadie *et al.* (LIGO Scientific Collaboration), *Astrophysical Journal* **722**, 1504 (2010), arXiv:1006.2535 [gr-qc].
- [33] Some lattices, e.g. A_n^* , are conveniently represented in m -dimensional space, by an $m \times m$ matrix \mathbf{G}' , of which the lattice occupies an n -dimensional subspace. Hence only n columns of \mathbf{G}' are linearly independent, and we take \mathbf{G} to be the $m \times n$ matrix comprising linearly independent columns of \mathbf{G}' .
- [34] The QL factorization may be calculated using the more commonly implemented QR factorization [e.g. 29] via $E(\mathbf{G}) = \mathbf{Q}'\mathbf{R}$, $E(\mathbf{Q}') = \mathbf{Q}$, $E(\mathbf{R}) = \mathbf{L}'$, where $E(\cdot)$ reverses the order of both the rows and columns of its argument.
- [35] R. McKilliam, I. Clarkson, W. Smith, and B. Quinn, in *International Symposium on Information Theory and Its Applications* (IEEE, 2008) pp. 1–5.
- [36] We note that line 6 of the algorithm listings in [35] is incorrect; it should read $i = n + 1 - \lfloor (n + 1)(z_t + 0.5) \rfloor$.
- [37] D. E. Knuth, *The Art of Computer Programming*, 2nd ed., Vol. 3 (Addison-Wesley, Upper Saddle River, NJ, 1998).
- [38] P. Jaranowski and A. Królak, *Living Reviews in Relativity* **15** (2012), 10.12942/lrr-2012-4, retrieved 2014-10-18.
- [39] D. Whitbeck, *Observational Consequences of Gravitational Wave Emission From Spinning Compact Sources*, Ph.D. thesis, The Pennsylvania State University (2006).
- [40] K. Wette *et al.*, *Classical and Quantum Gravity* **25**, 235011 (2008), arXiv:0802.3332 [gr-qc].
- [41] R. Prix, “Frequency metric for CW searches,” (2013), unpublished.
- [42] B. Abbott *et al.* (LIGO Scientific Collaboration), *Physical Review D* **76**, 082001 (2007), arXiv:gr-qc/0605028.
- [43] LIGO Scientific Collaboration and Virgo Collaboration, *The LSC-Virgo white paper on gravitational wave data analysis*, Tech. Rep. LIGO-T1300432-v2 (LIGO, 2013).
- [44] P. R. Brady and T. Creighton, *Physical Review D* **61**, 082001 (2000), arXiv:gr-qc/9812014.
- [45] C. Messenger, *Physical Review D* **84**, 083003 (2011), arXiv:1109.0501 [gr-qc].
- [46] P. Leaci and R. Prix, in preparation (2014).